

An argumentative agent-based model of scientific inquiry

AnneMarie Borg¹, Daniel Frey², Dunja Šešelja³, and Christian Straßer⁴

¹*Institute for Philosophy II, Ruhr-University Bochum*

²*Heidelberg University*

³*Institute for Philosophy II, Ruhr-University Bochum, Center for Logic and Philosophy of Science, Ghent University*

⁴*Institute for Philosophy II, Ruhr-University Bochum, Center for Logic and Philosophy of Science, Ghent University*

December 15, 2016

Abstract

In this paper we present an agent-based model (ABM) of scientific inquiry aimed at investigating how different social networks impact the efficiency of scientists in acquiring knowledge. As such, the ABM is a computational tool for tackling issues in the domain of scientific methodology and science policy. In contrast to existing ABMs of science, our model aims to represent the argumentative dynamics that underlies scientific practice. To this end we employ abstract argumentation theory as the core design feature of the model. This helps to avoid a number of problematic idealizations which are present in other ABMs of science and which impede their relevance for actual scientific practice.

1 Introduction

In this paper we present a computational tool for tackling issues in the domain of scientific methodology and science policy, which concern social aspects of scientific inquiry and the division of cognitive labor. Recent approaches to these questions have utilized agent-based models (ABMs) (e.g. [27, 28], [26], [4], [24], etc). One of the advantages of such formal approaches is that they avoid hasty generalizations resulting from the traditional case-study approach since they employ an experimental method in which a number of relevant variables can be controlled during the run of simulations. However, the main pitfall of ABMs is that they frequently suffer from a high degree of idealization and simplification,

which may sometimes impede the relevance of the results for the context of actual scientific inquiry [16, 24].

One such idealization concerns the question how we represent evaluations on the basis of which scientists choose paths to pursue. Ever since Laudan introduced the notion of the context of pursuit [13], many have argued that the assessment of one’s research direction has to be distinguished from the epistemic justification of one’s beliefs (e.g. [15, 21]). While in some ABMs these two types of evaluations have been conflated (e.g. [27, 28]), in others the former type of assessment has been represented in an oversimplified way (e.g. [26]). In both cases the way in which scientists respond to the information which they receive from their peers, and in view of which they update their pursuit-related attitudes, has not been modeled according to the idea of *heuristic appraisal* [15]. This idea assumes that when scientists face possible anomalies in their research programs, such anomalies trigger a search for counterarguments that could turn apparent refutations into confirmatory instances [12].

This points to another problematic idealization frequently present in ABMs, namely the fact that the interaction among scientists is represented as a simple update in view of received information. While some models have tried to relax this idealization by representing agents as “trusting” others only if they have sufficiently similar views (e.g. [10]), by introducing “noise” in the received information [4], by assigning different weights to the opinions of agents [18], or by assigning different epistemic systems to agents [3], these adjustments do not capture the argumentative nature of scientific interaction, such as the above mentioned search for a counterargument in face of an attack.

In this paper we offer a novel approach to the agent-based modeling of scientific inquiry, which is based on argumentation, and which aims to soften the above mentioned idealizations. In contrast to most other ABMs of science, our model is based on the idea that an essential component of scientific inquiry is an argumentative dynamics between scientists. To this end, we employ abstract argumentation frameworks as one of the design features of our ABM (previously shown fruitful for the modeling of scientific debates in [22] and employed in an ABM of social behavior in [7]).

The presented version of the model is designed to investigate how different social networks impact the efficiency of scientists in discovering the best of the pursued scientific theories. This question has previously been tackled in [27, 28] whose ABMs suggest that information sharing may sometimes be counter-productive for the efficiency in knowledge acquisition. This is due to the fact that scientists may initially get misleading results, suggesting that theory T_1 is better than theory T_2 , while from an objective point of view it is the other way around. If such misleading information is shared widely, the whole scientific community may start pursuing T_1 and abandon T_2 . However, recent discussion in [19] has shown that the results by Zollman [27, 28] regarding the harmful effects of increased interaction are not robust under some minor changes of the relevant parameters. Nevertheless, the authors also point out that similar results have been obtained by a structurally different model, namely the one by Grim [8], which leaves the question of the epistemic benefits of scientific inter-

action open. Moreover, both Zollman’s and Grim’s models have been criticized as suffering from various problematic idealizations in [20], which impede the relevance of their results for actual scientific practice. One of these simplifications is the above mentioned representation of scientific interaction as a simple update in view of new evidence. The aim of our model is to tackle the same question, while representing scientific interaction in a more adequate way.

The paper is structured as follows. In Section 2 we introduce the idea underlying our ABM and its main features. In Section 3 we present the central findings of the presented version of our model. In Section 4 we compare our findings with those obtained by other ABMs of science and discuss idealizations present in our current model, which should be taken into consideration when assessing the relevance of our results for the actual scientific practice. In Section 5 we suggest ideas for further enhancements of this ABM.

2 The model

The aim of our ABM is to represent scientists engaged in a scientific inquiry with the goal of finding the best of the given rivaling theories, where they occasionally exchange arguments with other scientists. In this section we will explain the central elements of our model, namely, the landscape which agents explore, the behavior of agents and the notion of social networks employed in our simulations.¹

2.1 The landscape

Agents, representing scientists, move along an *argumentative landscape*. The argumentative landscape, which represents rivaling theories in a given scientific domain, is based on a dynamic abstract argumentation framework. Let us explain what ‘abstract’ and ‘dynamic’ mean here.

An abstract argumentation framework (AF), introduced by Dung in [5] is a formal framework consisting of a set of abstract entities \mathcal{A} representing *arguments* and an *attack relation* \rightsquigarrow over \mathcal{A} . Similarly, the framework underlying our model consists of a set of arguments and an attack relation over this set. In addition to attacking each other, arguments may also be connected by a *discovery relation* \hookrightarrow . The latter represents the path which scientists have to take in order to discover different parts of the given theory, i.e., some argument b in the given theory can only be discovered after an argument a has been discovered for which $a \hookrightarrow b$.

A scientific theory is represented as a conflict-free set of arguments (i.e. no argument in the theory attacks an argument in the same theory), connected by discovery relations, resulting in a tree-like graph. Formally, an argumentative landscape is given by a triple $\langle \mathcal{A}, \rightsquigarrow, \hookrightarrow \rangle$ where $\mathcal{A} = \langle \mathcal{A}_1, \dots, \mathcal{A}_m \rangle$ is partitioned

¹The source code is available at <https://github.com/g4v4g4i/ArgABM/tree/AppArg2017-submission>.

in m many theories $T_i = \langle \mathcal{A}_i, a_i, \hookrightarrow \rangle$ which are trees with $a_i \in \mathcal{A}_i$ as a root and

$$\rightsquigarrow \subseteq \bigcup_{\substack{1 \leq i, j \leq m \\ i \neq j}} (\mathcal{A}_i \times \mathcal{A}_j) \quad \text{and} \quad \hookrightarrow \subseteq \bigcup_{1 \leq i \leq m} (\mathcal{A}_i \times \mathcal{A}_i).$$

Specifying \rightsquigarrow like this ensures that the theories are conflict-free.

The abstractness of the framework concerns all of its elements. Instead of providing the concrete content and structure of the given arguments, we represent them as abstract entities. Similarly, we do not reveal the concrete nature of the attack or the discovery relation.

The framework is dynamic in the sense that agents gradually discover arguments, as well as attack and discovery relations between them. Given the abstract nature of arguments, we interpret them as hypotheses which scientists investigate, occasionally encountering defeating evidence, represented by attacks from other arguments, and then attempting to find defending arguments for the attacked hypothesis a (i.e., to find arguments in the same theory attacking arguments from other theories that attack a). This dynamic aspect is implemented by associating arguments with their *degree of exploration* for an agent at a given time point of a run of the simulation: for each agent \mathbf{ag} and each argument $a \in \mathcal{A}$, $\mathbf{expl}(a, \mathbf{ag}) \in \{0, \dots, 6\}$ where 0 indicates that the argument is unknown to \mathbf{ag} and 6 indicates that the argument is fully explored and cannot be further explored.² In view of this agents have subjective and limited insights into the structure of the landscape. Whether an attack or discovery relation between two arguments a and a' is visible to an agent \mathbf{ag} depends on the degree of exploration $\mathbf{expl}(a, \mathbf{ag})$: the higher $\mathbf{expl}(a, \mathbf{ag})$ is, the more relation[s] between a and other arguments will be visible (additionally agents may learn about the landscape by communicating with other agents, see Section 2.3).

2.2 Basic behavior of agents

The model is round-based and each round agents perform actions which are among the following:

1. exploring a single argument, thereby gradually discovering possible attacks (on it, and from it to an argument that belongs to another theory) as well as discovery relations to neighboring arguments;
2. moving to a neighboring argument along the discovery relation within the same theory;
3. moving to an argument of a rivaling theory.

While agents start the run of the simulation at the root of a given theory, they will gradually discover more and more of the argumentative landscape. This way each turn an agent operates on her own (subjective) fragment of the

²Our model is round-based (more on that in Section 2.2). Each round may be interpreted as one research day. Since each of the 6 levels of an argument is explored in 5 rounds, each argument represents a hypothesis that needs 30 research days to be fully investigated.

landscape, which consists of her discovered arguments which are explored by her to a specific degree, and her discovered (attack and discovery) relations between the arguments.

In order to decide whether to work on the current theory (items 1 and 2 above), or whether to better start working on an alternative theory (item 3) agents are equipped with the ability to evaluate theories. Every few rounds agents apply an evaluative procedure, which is based on the *degree of defensibility* of each of the given theories. A theory has degree of defensibility n if it has n defended arguments where an argument a is defended in the theory if each attacker b from another theory is itself attacked by some argument c in the current theory. Agents decide to move to a rivaling theory if the degree of defensibility of their current theory is below a relative threshold compared to the theory with the highest degree of defensibility, i.e., the theory with the most defended arguments.

To make a decision between options 1 and 2, each agent employs the following heuristic: at every time step she considers all arguments in her direct neighborhood (relative to the discovery relation) that could possibly be the next ones to work on. With a certain probability she will then move to one of these arguments, or alternatively keep on exploring her current argument. In case the argument she's currently at is fully explored, she will try to move on to a next neighboring argument, and if such an argument isn't visible (e.g., if she has reached the end of a branch of her theory) she will try to move to the parent argument, or in case it is fully explored, she will move to another not fully explored argument in the same theory.

The decision making of agents also includes some prospective considerations. If during her exploration an agent discovers an attack on the argument a she is currently investigating, she will attempt to discover a defense for it. A defending argument may be found among the visible neighboring arguments of a (relative to the discovery relation). In case a is attacked by arguments b_1, \dots, b_n , an agent **ag** working on a will 'see' outgoing attack arrows $a' \rightsquigarrow b_i$ (where $1 \leq i \leq n$) from an already discovered child argument a' of a , even in cases where $a' \rightsquigarrow b_i$ is not yet discovered by **ag** since a' is insufficiently explored by **ag**.³ This way our agent knows that exploring a' may help in defending a . If no potential defender of a is visible she keeps on exploring a in the hope of discovering new neighbors and thus new potential defenders.

This idea corresponds to the situation in which a scientist discovers a problem in her hypothesis and attempts to find a way to resolve it. While she may not have a solution ready at hand, she may have a method for finding such a solution (for example, going back to the laboratory and conducting some new experiments).⁴

³In such cases, where an attack relation is merely 'seen' but not yet discovered by an agent, it is not yet considered as a defense when the respective theory is evaluated by the agent.

⁴Such a heuristic response to apparent refutations belongs to what Lakatos has dubbed the *negative heuristics* of a research program [12].

2.3 Social networks

An agent discovers the argumentative landscape by investigating arguments (as described in Section 2.2) or by means of exchanging information about the landscape with other agents. We will now discuss the later aspect. As mentioned above, the presented version of the model is designed to investigate how different social networks impact the efficiency of scientists in discovering the best theory. In contrast to other ABMs employing the idea of social networks (e.g. [27, 28, 8]), which represent connectivity only in view of different types of graphs that connect agents, we distinguish between two types of social networks.

First, our agents are divided into *collaborative networks* that may consist of individuals working on the same theory (homogeneous groups), or of individuals working on different theories (heterogeneous groups). While each agent gathers information (i.e. the attack and discovery relations between arguments) on her own, every five steps this information is shared with all other agents forming the same collaborative network.

Second, besides sharing information with agents from the same network, every five steps each agent shares information with agents from other collaborative networks with a given *probability of information sharing* that is determined before the run of the simulation.⁵ This way the agents form ad-hoc and random networks with agents from other research collaborations. A higher probability of information sharing leads to a higher degree of interaction among agents.

Finally, we represent *reliable* and *biased* scientists by allowing for different approaches to the sharing of information between networks. A reliable agent shares all the information she has gathered during her exploration of the current theory, while a biased agent does not share the information regarding the discovered attacks on her current theory.

Agents share information either in a unidirectional (an agent sends information to another agent) or a bidirectional way (agents exchange information one with another). Moreover, our model takes into account the fact that receiving information is time costly: when an agent receives information, she will not explore the argument on which she is standing nor move. This corresponds to the idea that scientists need to invest time when reading papers by other scientists, which they would otherwise devote to their own research.

3 The main findings

In this section we will first specify the parameters used in the simulations and then present our main results.

⁵While agents share their full subjective knowledge within their respective collaborative networks, the information which they share with agents from other networks concerns recently obtained knowledge of the theory which they are currently exploring. This corresponds to a situation in which a scientist writes a paper that presents arguments for and/or against hypotheses regarding the theory she is currently pursuing.

3.1 Parameters used in simulations

We have run the simulation 100 times with 10, 20, 30, 40, 70 and 100 agents by varying the following settings:

1. different probabilities of an agent communicating with agents from other collaborative networks, namely: 0, 0.3, 0.5, and 1;
2. different types of collaborative networks, namely: homogeneous and heterogeneous ones;
3. different approaches to communicating, namely: reliable and biased agents;
4. two different landscapes, one representing two theories, and one representing three theories.⁶

The program runs until each agent is on a fully explored theory, since at this point the agents are locked in the given theories and no further information is available, in view of which they would move to another theory. Our main research question is how efficient agents are in each of the above listed scenarios, where efficiency is assessed in terms of the success of agents in acquiring knowledge, as well as in terms of the time needed for the run to be completed.

In this version of the model we have defined success in the following way: if, at the end of the run, there is no theory for which the number of agents working on it is greater than the number of agents working on the objectively best theory, the run is considered successful. In contrast to some other ABMs of science, which define success in terms of convergence of all agents onto the best theory (e.g. in [27, 28]), our notion is obviously weaker. Our choice was motivated by a pluralist view on scientific inquiry, according to which, a parallel existence of rivaling scientific theories is epistemically and heuristically beneficial for the goals of a scientific community (e.g. [14, 11, 2]). A primary epistemic concern is thus not the convergence of all scientists onto the same theory, but rather assuring that the best theory is among the most actively investigated ones.

3.2 Preliminary results

In what follows we present the most significant results of our simulations. For each of the comparisons below we will vary only the parameter given in the respective paragraph title while keeping all other parameters fixed.

Increased information sharing. Increased information sharing appears to be beneficial for both reliable and biased groups. On the one hand, homogeneous groups that share information with probability 1 perform faster, while being similarly successful as those with lower strict positive probabilities (SPP) of information sharing (Fig. 1 and 3). As expected, homogeneous groups that do not share information among each other perform the worst in terms of both,

⁶Some other parameters are as follows. Theories are modeled as trees of depth 3, where each argument (except for the final leaves) has 4 child-arguments. The extent to which each of the theories is attacked is specified in the interface of the model in terms of the probability that each of the arguments of the given theory is attacked. Here we have opted for 0.3 probability since we wish to represent a situation in which rivaling theories are not completely problematic (as would be e.g. pseudo-scientific theories).




Figure 1: Success, reliable, homogeneous groups

success and time. On the other hand, heterogeneous groups that share information are all similar for SPP in terms of time and success, (see Fig. 2 and 3), except for larger groups (70 and 100 agents) which are faster in case the probability of sharing information is 1. Groups with probability 0 perform worse in both respects.

Reliable vs. biased agents. With respect to homogeneous groups, reliable agents appear to be more successful than the biased ones in case of SPP of information sharing, while sometimes being only slightly slower. With respect to heterogeneous groups, reliable and biased groups perform similarly in terms of both success and time.

Homogeneous vs. heterogeneous networks. In case of 0.3 and 0.5 probabilities of information sharing, heterogeneous networks tend to be faster than the homogeneous ones, while being similar in terms of success. In case the probability is 1, heterogeneous and homogeneous groups are similar in both respects (Fig. 3). In case of no information sharing heterogeneous groups are much more successful. These results hold for both reliable and biased agents.

4 Discussion

In this section we will first compare our results with those obtained by other ABMs of science, and then we will turn to a critical analysis of some idealizations present in our model.




Figure 2: Success, reliable, heterogeneous groups




Figure 3: Time needed

Our finding that increased communication tends to be epistemically beneficial (or at least, not epistemically harmful) undermines the robustness of conclusions drawn from ABMs in [27, 28, 8, 9], under different modeling choices. As we argue below, there is no reason to assume that any of these ABMs represent scientific interaction more adequately than the presented ABM does. In addition, our finding that biased agents perform worse than reliable ones under conditions of an increased information flow shows that once we apply the notion of bias to the way in which agents share information (rather than to their confidence in the given theories with which they begin an inquiry, as done in [28]), we get results that are contrary to [28].

While a number of open issues regarding the comparison of our results with those obtained by means of other ABMs remain for future research, we can already highlight the features of our model that allow for a more adequate representation of scientific inquiry than this has so far been done with ABMs. First of all, our model addresses problematic features identified in other ABMs (see Section 1), namely (i) the inadequate representation of heuristic appraisal, as well as (ii) the inadequate representation of information flow among scientists. Regarding (i), scientists in our model are equipped with certain prospective considerations, explained in Section 2.2. In the current version of the model, such prospective considerations play a role in methods that guide agents in their inquiry. It remains a task for a future version of the model to also incorporate such features into the evaluations performed by agents in view of which they can judge how promising their theory is. Such an assessment would more aptly capture heuristic appraisal, which informs scientists how worthy of pursuit different theories are, rather than how confirmed (or defensible) they are in view of the available evidence.

Regarding (ii), the notion of information flow is in contrast to many other ABMs of science not just a simple update of information. Instead, exchange of information is represented as argumentative, which means that received information is critically assessed. Hence, when an agent receives new information regarding a discovered argument, she will assess it as acceptable (in case it can be defended in view of her knowledge base) or as unacceptable (in case it cannot be defended in view of her knowledge base).

Moreover, different information triggers different heuristic activities on the side of the receiver, so that an agent who has discovered an attack on her current argument tries to find its defense. Finally, we acknowledge the fact that receiving information (such as reading articles by other scientists) may be time costly.

Nevertheless, our ABM is still based on a number of idealizations, the impact of which should be examined in future research. First, heuristics of agents are highly simplified, including their search for defense in view of discovered attacks. Even though agents may recognize a defense in case it is located in their surroundings (i.e. in one of the child-arguments of the given attacked argument, which an agent currently explores), a defense may be present at an entirely different branch of the tree. Equipping agents with more insight into where a defense may be, and thus representing a scientist as having methods

for finding solutions for the current anomalies of the theory, is another task for future research.

Second, anomalies of a given theory are currently represented solely as attacks from one of the rivaling theories. An improved version of the model should allow for anomalies to be represented as counter-evidence discovered only as a result of exploring the given theory. Future versions of our model will include a more direct representation of evidence.

Third, what parameter settings for building landscapes in our model are representative of specific types of scientific controversies is an open question.

In view of these remarks, it is important to interpret the results of our runs cautiously when it comes to their relevance for the actual scientific practice. Just like all other existing ABMs of science, our model is still just a “bookshelf model”, which means that it is still too simplified to be fully informative of actual scientific practices.⁷ Nevertheless, as we have argued, the highly modular nature of our model together with its specific design features makes it significantly closer to the aim of representing scientific inquiry than the currently existing ABMs. As such, our model offers a fruitful basis for further improvements, which can provide insights into real world phenomena.

Finally, let us compare our model with Gabbriellini and Torroni’s (G&T) ABM [7]. Their aim is to study polarization effects among communicating agents, for instance, in online debates. Similarly to our approach, their model is based on an abstract argumentation framework. Agents start with an individual partial knowledge of the given framework and enhance their knowledge by means of communication. Since G&T do not model inquiry, their agents cannot discover new parts of the graph by means of ‘investigating’ arguments. Rather, they exchange information by engaging in a dialogue modeled after Mercier & Sperber’s argumentative theory of reasoning [6]. This way, agents may learn about new arguments and attacks but also remove attack relations. Whether new information is incorporated in the knowledge of an agent depends on the trust relation between the discussants. The beliefs of agents are represented by applying Dung-style admissibility-based semantics to the known part of the argumentation framework of an agent. This is quite different from our model where the underlying graph topology is given by several discovery trees of arguments representing scientific theories and attacks between them. This additional structure of the argumentation graph is essential since we do not model the agents’ beliefs in individual arguments but rather evaluative stances of agents that inform their practical decision of which theory to work on. While an admissibility-based semantics would lead to extensions that feature unproblematic sets of arguments from different theories (ones that form conflict-free and fully defensible sets), in our approach agents pick theories to work on. For this, they compare the merits of the given theories, pick the one that is ‘most’ defended (where typically no theory is fully defended before the end of a run), and employ heuristic behavior to tackle open problems of theories (represented

⁷The importance of distinguishing between “bookshelf models” and those that are relevant for real world phenomena has been emphasized in [17] in the context of economic models. The same considerations apply to models of scientific inquiry.

by incoming attacks). It will be the topic of future research to include dialogue protocols that are relevant for scientific communication, such as information-seeking, inquiry and deliberation dialogues [25].

5 Outlook and conclusion

In this paper we have presented an ABM of scientific inquiry which makes use of abstract argumentation, aiming to model the argumentative nature of scientific inquiry. The presented version of our model is designed to tackle the question how different degrees of information flow among scientists affects the efficiency of their knowledge acquisition. Our results suggest that an increased information sharing is epistemically beneficial, which undermines the robustness of contrary results obtained by previous ABMs of science, under different modeling choices. While we have argued that our model represents scientific inquiry more adequately than the previous ABMs of science, we have also emphasized a number of issues that remain to be tackled in future research. We will conclude the paper by showing the fruitfulness of our ABM for the investigation of related questions concerning social aspects of scientific inquiry.

First, our model can be enhanced with different types of research behavior, such as “mavericks” and “followers”, introduced in [26]. Next, the model can be enhanced with other aspects of scientific inquiry, such as an explanatory relation and a set of explananda [22]. This would allow for an investigation of different evaluative procedures which agents perform when selecting their preferred theory (e.g. in addition to the degree of defensibility, agents can take into account how much their current theory explains, or how well it is supported by evidence). Furthermore, a number of enhancements available from the literature on AFs, such as probabilistic semantics [23], values [1], etc. can be introduced in future versions of our ABM.

Acknowledgements The research of Annemarie Borg and Christian Straßer was supported by a Sofja Kovalevskaja award of the Alexander von Humboldt-Foundation, funded by the German Ministry for Education and Research.

References

- [1] Bench-Capon, T.J.M.: Value based argumentation frameworks. arXiv cs.AI/0207059 (2002)
- [2] Chang, H.: Is Water H₂O? Evidence, Pluralism and Realism. Springer (2012)
- [3] De Langhe, R.: Peer disagreement under multiple epistemic systems. Synthese 190, 2547–2556 (2013)
- [4] Douven, I.: Simulating peer disagreements. Studies in History and Philosophy of Science Part A 41(2), 148–157 (2010)

- [5] Dung, P.M.: An argumentation-theoretic foundation for logic programming. *The Journal of logic programming* 22(2), 151–171 (1995)
- [6] Gabbriellini, S., Torroni, P.: MS dialogues: Persuading and getting persuaded. a model of social network debates that reconciles arguments and trust. *Proc. 10th ArgMAS* (2013)
- [7] Gabbriellini, S., Torroni, P.: A new framework for ABMs based on argumentative reasoning. In: *Advances in Social Simulation*, pp. 25–36. Springer (2014)
- [8] Grim, P.: Threshold phenomena in epistemic networks. In: *AAAI Fall Symposium: Complex Adaptive Systems and the Threshold Effect*. pp. 53–60 (2009)
- [9] Grim, P., Singer, D.J., Fisher, S., Bramson, A., Berger, W.J., Reade, C., Flocken, C., Sales, A.: Scientific networks on data landscapes: question difficulty, epistemic success, and convergence. *Episteme* 10(04), 441–464 (2013)
- [10] Hegselmann, R., Krause, U.: Opinion dynamics driven by various ways of averaging. *Computational Economics* 25(4), 381–405 (2005)
- [11] Kitcher, P.: *Science in a democratic society*. Prometheus Books (2011)
- [12] Lakatos, I.: *The methodology of scientific research programmes*. Cambridge University Press, Cambridge (1978)
- [13] Laudan, L.: *Progress and its Problems: Towards a Theory of Scientific Growth*. Routledge & Kegan Paul Ltd, London (1977)
- [14] Longino, H.: *The Fate of Knowledge*. Princeton University Press, Princeton (2002)
- [15] Nickles, T.: Heuristic appraisal: Context of discovery or justification? In: Schickore, J., Steinle, F. (eds.) *Revisiting Discovery and Justification: Historical and philosophical perspectives on the context distinction*, pp. 159–182. Springer, Netherlands (2006)
- [16] Payette, N.: Agent-based models of science. In: Scharnhorst, A., Börner, K., van den Besselaar, P. (eds.) *Models of Science Dynamics*, pp. 127–157. *Understanding Complex Systems*, Springer (2012)
- [17] Pfeleiderer, P.: Chameleons: The misuse of theoretical models in finance and economics. *Revista de Economía Institucional* 16(31), 23–60 (2014)
- [18] Riegler, A., Douven, I.: Extending the hegselmann–krause model iii: From single beliefs to complex belief states. *Episteme* 6(02), 145–163 (2009)
- [19] Rosenstock, S., O’Connor, C., Bruner, J.: In epistemic networks, is less really more? *Philosophy of Science* (2016)

- [20] Šešelja, D.: What have we learned from agent based models about the epistemic effects of interaction among scientists? <http://tinyurl.com/jrzcp7> (2017)
- [21] Šešelja, D., Kosolosky, L., Straßer, C.: Rationality of scientific reasoning in the context of pursuit: drawing appropriate distinctions. *Philosophica* 86, 51–82 (2012)
- [22] Šešelja, D., Straßer, C.: Abstract argumentation and explanation applied to scientific debates. *Synthese* 190, 2195–2217 (2013)
- [23] Thimm, M.: A probabilistic semantics for abstract argumentation. In: *ECAI*. pp. 750–755 (2012)
- [24] Thoma, J.: The epistemic division of labor revisited. *Philosophy of Science* 82(3), 454–472 (2015)
- [25] Walton, D., Krabbe, E.C.: *Commitment in dialogue: Basic concepts of interpersonal reasoning*. SUNY press (1995)
- [26] Weisberg, M., Muldoon, R.: Epistemic landscapes and the division of cognitive labor. *Philosophy of science* 76(2), 225–252 (2009)
- [27] Zollman, K.J.S.: The communication structure of epistemic communities. *Philosophy of Science* 74(5), 574–587 (2007)
- [28] Zollman, K.J.S.: The epistemic benefit of transient diversity. *Erkenntnis* 72(1), 17–35 (2010)